

CLAIMS

- 1 1. A method for selecting a coredump disk from a set of disks owned by a
2 failed server into which memory data (“coredump”) of the failed server is stored, the
3 method comprising the steps of:
- 4 (a) identifying available disks that can receive a coredump;
- 5 (b) ordering the disks identified in step (a) so as to prefer disks that are least
6 likely to be needed for normal service by the server;
- 7 (c) further ordering the disks so as to prefer disks that require a least amount of
8 preparation to receive a coredump; and
- 9 (d) selecting as a best candidate for use as the coredump disk a first one of the
10 disks in the order after steps (a), (b), (c) and (d).
- 1 2. The method as set forth in claim 1 further comprising, determining
2 whether the selected coredump disk is formatted for use in normal service and if format-
3 ted, writing a “not-formatted” attribute to labels of the coredump disk.
- 1 3. The method as set forth in claim 1 further comprising, maintaining an at-
2 tribute in a core region header of the coredump disk to indicate the current status of the
3 coredump, the status including each of “complete coredump,” “no coredump,”
4 “coredump in progress” and “coredump aborted.”
- 1 4. The method as set forth in claim 3 further comprising, writing a
2 “coredump in progress” signature to the coredump disk so as to prevent the coredump
3 disk from being used for normal file service including updating the coredump status at-
4 tribute to indicate the status of “coredump in-progress” and writing the updated
5 “coredump in-progress” coredump status attribute to the core region header of the
6 coredump disk.

1 5. The method as set forth in claim 1 further comprising providing pointers
2 in the core region header that point to a file system region of the coredump disk in which
3 the coredump is stored.

1 6. The method as set forth in claim 5 further comprising, writing the
2 coredump to the file system region.

1 7. The method as set forth in claim 1 further comprising, where a spare disk
2 is unavailable for selection as the coredump disk, providing coredump data to a core re-
3 gion in each of the set of disks in a distributed manner and, after one of either the re-
4 booting by the failed filer or taking-over ownership of the set of disks by the taking-over
5 ownership other filer, formatting and writing the data of the coredump to a root file sys-
6 tem of either of the rebooted failed filer or the taking-over ownership other filer, respec-
7 tively.

1 8. A method of identifying a coredump disk, during one of either rebooting
2 by the failed filer or taking-over ownership the set of disks by another filer, from a set of
3 disks owned by a failed filer into which memory data ("coredump") of the failed filer is
4 stored, the method comprising the steps of:

5 locating the coredump disk by reading labels on the disks of the set of disks to lo-
6 cate any of the disks that are non-formatted spare disks;

7 caching the labels for later use by a disk label assimilation process; and

8 reading a core region header in a respective label of the cached labels of each of
9 the non-formatted spare disks to locate a coredump status attribute in the respective of the
10 labels.

1 9. The method as set forth in claim 8 further comprising, freeing the cached
2 labels after either one of the rebooting or taking-over ownership occurs and the disk label
3 assimilation process completes.

1 10. The method as set forth in claim 8 further comprising, detecting a com-
2 pleted coredump in the coredump disk by locating a coredump status attribute in the
3 coredump disk and thereafter formatting and writing data of the coredump to a root file
4 system of either the rebooted failed filer or the taking-over ownership other filer.

1 11. The method as set forth in claim 10 further comprising, returning the
2 coredump disk to a "hot" spare status from a status in which the coredump disk is un-
3 available for use in normal file service including updating the coredump status attribute
4 to indicate a status of "no coredump" and writing the updated "no coredump" coredump
5 status attribute data to the core region header of the coredump disk.

1 12. The method as set forth in claim 10 further comprising, writing a
2 "coredump in progress" signature to the coredump disk so as to prevent the coredump
3 disk from being used for normal file service including updating the coredump status at-
4 tribute to indicate a status of "coredump in-progress" and writing the updated "coredump
5 in-progress" coredump status attribute to the core region header of the coredump disk.

1 13. The method as set forth in claim 8 further comprising, where a spare disk
2 is unavailable for selection as the coredump disk, providing coredump data to a core re-
3 gion in each of the set of disks in a distributed manner and, after one of either the re-
4 booting by the failed filer or taking-over ownership of the set of disks by the taking-over
5 ownership other filer, formatting and writing the data of the coredump to a root file sys-
6 tem of either of the rebooted failed filer or the taking-over ownership other filer, respec-
7 tively.

1 14. A method returning a coredump disk, selected from a set of disks owned
2 by a failed filer into which memory data ("coredump") of the failed filer is stored, back to
3 a "hot" spare status comprising the steps of:
4 recognizing an event that requires a spare disk to be made available; and

5 writing a "kill" signature to a core region header of the coredump disk in order to
6 return the coredump disk as soon as practicable to the "hot" spare status.

1 15. The method as set forth in claim 14 wherein the step of writing includes
2 updating a coredump status attribute in the core region header to indicate a "no
3 coredump" status from one of at least of "coredump in-progress" status or a "coredump
4 complete" status and writing the updated core region header to the coredump disk.

1 16. A disk mapping layout for a spare disk that is one of a set of disks owned
2 by a failed file server into which memory data ("coredump") of the failed file server is
3 stored comprising:

4 a disk table of contents (TOC) region including a disk TOC that maps a plurality
5 of regions on the disk;

6 a core region that stores an identifier indicating a status of the disk including at
7 least one of a no-coredump status, a coredump in-progress status, a coredump complete
8 status and a coredump aborted status; and

9 a file system region including storage space for the pointed-to by the core region
10 that is adapted to store a complete coredump data therein so as to define the disk as a
11 dedicated coredump disk.

1 17. The disk mapping layout as set forth in claim 16 further comprising a pair
2 of disk labels regions located at a fixed offset and remote from each other so as to mini-
3 mize risk of loss of labels data.

1 18. A method for selecting from a pool of candidate data storage devices
2 owned by a server that includes a set of data storage devices with which normal file
3 service by the server is performed, a best candidate data storage device for receiving in a
4 predetermined coredump storage space from the server memory contents (a "coredump")
5 comprising the steps of:

6 identifying data storage devices in the pool of candidate data storage devices that
7 are available and sized to receive the coredump; and

8 selecting the best candidate data storage device based upon a predetermined com-
9 bination criteria including: (i) which data storage device of the pool of candidate data
10 storage devices is adapted to complete reception of the coredump in the shortest time; (ii)
11 which data storage device of the pool of candidate data storage devices is adapted to
12 complete reception of the coredump with the least disruption of the normal file service;
13 and (iii) which data storage device of the pool of candidate data storage devices is
14 adapted to receive the coredump with minimal excess of storage space in the predeter-
15 mined coredump storage region after reception is complete.

1 19. The method as set forth in claim 18 wherein the step of identifying further
2 comprises selecting any available spare data storage device from the pool of candidate
3 storage devices.

1 20. The method as set forth in claim 19 wherein the step of selecting further
2 comprises:

3 (a) selecting from the identified spare data storage devices those having a core
4 region, including a core region header;

5 (b) further selecting from the identified spare data storage devices in step (a)
6 those having a file system region of sufficient space to receive the coredump; and

7 (c) ordering the identified spare data storage devices selected by steps (a) and
8 (b) so as to prefer data storage devices least likely to be needed for normal file service by
9 the file server;

10 (d) further ordering the identified spare data storage devices so as to prefer the
11 disks requiring a least amount of preparation to receive the coredump; and

12 (e) selecting as the best candidate, the one of the identified data storage de-
13 vices, which is the first data storage device yielded after the steps (a), (b), (c) and (d).

1 21. A computer-readable medium including program instructions for selecting
2 from a pool of candidate data storage devices owned by a server that includes a set of
3 data storage devices with which normal file service by the server is performed, a best
4 candidate data storage device for receiving in a predetermined coredump storage space
5 from the server memory contents (a “coredump”), the program instructions performing
6 the steps of:

7 identifying data storage devices in the pool of candidate data storage devices that
8 are available and sized to receive the coredump; and

9 selecting the best candidate data storage device based upon a predetermined com-
10 bination criteria including: (i) which data storage device of the pool of candidate data
11 storage devices is adapted to complete reception of the coredump in the shortest time; (ii)
12 which data storage device of the pool of candidate data storage devices is adapted to
13 complete reception of the coredump with the least disruption of the normal file service;
14 and (iii) which data storage device of the pool of candidate data storage devices is
15 adapted to receive the coredump with minimal excess of storage space in the predeter-
16 mined coredump storage region after reception is complete.

1 22. The computer-readable medium as set forth in claim 21 wherein the step
2 of identifying further comprises selecting any available spare data storage device from
3 the pool of candidate storage devices.

1 23. The computer-readable medium as set forth in claim 21 wherein the step
2 of selecting further comprises:

3 (a) selecting from the identified spare data storage devices those having a core
4 region, including a core region header;

5 (b) further selecting from the identified spare data storage devices in step (a)
6 those having a file system region of sufficient space to receive the coredump; and

7 (c) ordering the identified spare data storage devices selected by steps (a) and
8 (b) so as to prefer data storage devices least likely to be needed for normal file service by
9 the file server;

10 (d) further ordering the identified spare data storage devices so as to prefer the
11 disks requiring a least amount of preparation to receive the coredump; and

12 (e) selecting as the best candidate, the one of the identified data storage de-
13 vices, which is the first data storage device yielded after the steps (a), (b), (c) and (d).

1 24. A system for selecting from a pool of candidate data storage devices
2 owned by a server that includes a set of data storage devices with which normal file
3 service by the server is performed, a best candidate data storage device for receiving in a
4 predetermined coredump storage space from the server memory contents (a "coredump")
5 comprising the steps of:

6 an identifier that identifies data storage devices in the pool of candidate data stor-
7 age devices that are available and sized to receive the coredump; and

8 a selector that selects the best candidate data storage device based upon a prede-
9 termined combination criteria including: (i) which data storage device of the pool of can-
10 didate data storage devices is adapted to complete reception of the coredump in the short-
11 est time; (ii) which data storage device of the pool of candidate data storage devices is
12 adapted to complete reception of the coredump with the least disruption of the normal file
13 service; and (iii) which data storage device of the pool of candidate data storage devices
14 is adapted to receive the coredump with minimal excess of storage space in the predeter-
15 mined coredump storage region after reception is complete.

1 25. The system as set forth in 24 wherein the identifier is constructed and ar-
2 ranged to select any available spare data storage device from the pool of candidate stor-
3 age devices.

1 26. The system as set forth in claim 25 wherein the selector is constructed and
2 arranged to:

3 (a) select from the identified spare data storage devices those having a core
4 region, including a core region header;

- 5 (b) further select from the identified spare data storage devices in step (a)
6 those having a file system region of sufficient space to receive the coredump; and
7 (c) order the identified spare data storage devices selected by steps (a) and (b)
8 so as to prefer data storage devices least likely to be needed for normal file service by the
9 file server;
10 (d) further order the identified spare data storage devices so as to prefer the
11 disks requiring a least amount of preparation to receive the coredump; and
12 (e) select as the best candidate, the one of the identified data storage devices,
13 which is the first data storage device yielded after the selector performs (a), (b), (c) and
14 (d).

1 27. A method releasing a coredump disk, selected from a set of disks owned
2 by a failed filer into which memory data (“coredump”) of the failed filer is stored, back to
3 a status as a spare disk comprising the steps of:
4 maintaining the coredump disk as a discrete dedicated disk for receiving the
5 coredump and not available for normal file service by the failed filer; and
6 releasing the coredump disk from being the dedicated disk in response to at least
7 one of either (a) completion of the receipt of the coredump and transfer of the coredump
8 data to a predetermined data structure for subsequent access or (b) abort of the coredump
9 before the completion and transfer.

1 28. The method as set forth in claim 27 wherein the step of releasing includes
2 writing an attribute to a header of the coredump disk that is identified by either of the
3 failed filer or another filer taking-over ownership of the set of disks as no longer indicat-
4 ing a coredump status.

1 29. A utility for retrieving coredump data a coredump disk, selected from a set
2 of disks owned by a failed filer into which memory data (“coredump”) of the failed filer
3 is stored, comprising:

4 a reader that scans disk labels of the set of disks to locate a label indicating the
5 coredump disk with coredump being present in a file system region thereof; and
6 a writer that writes the coredump in file system region of the coredump disk to a
7 root file system of one of either the failed filer or another filer taking-over ownership of
8 the set of disks.